

A First Look into Transnational Routing Detours

Anne Edmundson, Roya Ensafi, Nick Feamster, Jennifer Rexford
Princeton University

Abstract

An increasing number of countries are passing laws that facilitate the mass surveillance of their citizens. In response, governments and citizens are increasingly paying attention to the countries that their Internet traffic traverses. In some cases, countries are taking extreme steps, such as building new IXPs and encouraging local interconnection to keep local traffic local. We find that although many of these efforts are extensive, they are often futile, due to the inherent lack of hosting and route diversity for many popular sites. We investigate how the use of overlay network relays and the DNS open resolver infrastructure can prevent traffic from traversing certain jurisdictions.

Categories and Subject Descriptors: C.2.1 [Computer-Communication Networks] *Network Architecture and Design*

General Terms: Measurement; Security

Keywords: surveillance; routing

1 Introduction

When Internet traffic enters a country, it becomes subject to those countries' laws. As a result, users have more need than ever to determine—and control—which countries their traffic is traversing. As an increasing number of countries pass laws that facilitate mass surveillance of their citizens [2], governments and citizens are increasingly motivated to divert their Internet traffic from countries that perform surveillance (notably, the United States [1]).

Many countries—notably, Brazil—are taking impressive measures to reduce the likelihood that Internet traffic transits the United States [1] including building a 3,500 mile long fiber-optic cable from Fortaleza to Portugal (with no use of American vendors); pressing companies such as Google, Facebook, and Twitter (among others) to store data locally; and switching its dominant email system (Microsoft Outlook) to a state-developed system called Expresso [1]. Brazil is

also building Internet Exchange Points (IXPs), now has the largest national ecosystem of public IXPs in the world, and the number of internationally connected ASes continues to grow. Brazil is not alone: IXPs are proliferating in eastern Europe, Africa, and other regions, in part out of a desire to “keep local traffic local”. Building IXPs alone, of course, cannot guarantee that Internet traffic for some service does not enter or transit a particular country: Internet protocols have no notion of national borders, and interdomain paths depend in large part on existing interconnection business relationships (or lack thereof). In this poster, we present evidence to suggest that existing Internet hosting and interdomain paths still make it difficult to avoid certain countries for many popular websites.

Prior work explored how central different countries are to interdomain routing based on simulated paths and an Internet topology [4]. Our study differs by actively measuring and analyzing the traffic originating in five different countries: Brazil, Netherlands, Kenya, India, and the United States. Using RIPE Atlas probes and the MaxMind geolocation service, we measure the country-level traffic paths for the Alexa Top 100 domains in each respective country. Using the current state of routing as a baseline for comparison, we then measure how avoidable a given country is to a client in either Brazil, Netherlands, India, Kenya, or the United States, using open resolvers and using the overlay network. Our contributions include:

- The first in-depth measurement study of nation-state routing for Brazil, Netherlands, Kenya, India, and the United States.
- A preliminary evaluation of how open DNS resolvers and overlay network relays can help citizens and governments discover and use paths that avoid certain countries.

We find that hosting for many popular websites lacks diversity; in many cases, even websites that are popular *locally* are hosted outside the country where citizens are trying to access them. For example, more than 50% of the Alexa Top 100 domains in Brazil are hosted in the United States. Internet paths also lack geographic diversity: About half of the paths originating in Kenya to the most popular Kenyan websites traverse the United States or Great Britain. Much of this phenomenon is due to “tromboning”, whereby an Internet path starts and ends in a country, yet transits an intermediate country; for example, about 13% of the paths that we explored from RIPE Atlas nodes in Brazil to the Alexa Top 100 in Brazil trombone through the United States. Fortunately, our

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

SIGCOMM'16, August 22–26, 2016, Florianópolis, Brazil.

Copyright 2016 ACM. ISBN 978-1-4503-4193-6/16/08...\$15.00

DOI: <https://dx.doi.org/10.1145/2934872.2959081>

preliminary results suggest that the use of overlay network relays to intentionally introduce network detours, and the use of open DNS resolvers to discover hosting diversity can reduce tromboning and generally help users select paths that avoid certain countries.

2 Method

Our measurement study tackles two questions: (1) Which countries do *default* Internet routing paths traverse?; (2) What types of methods can we use to take advantage of hosting and path diversity to help governments and citizens better control transnational Internet paths?

The first step of our measurement method uses RIPE Atlas probes in the country of interest to locally resolve each domain in the Alexa Top 100 and run traceroutes to the IP addresses in the DNS responses. The measurements were run using Paris traceroute and each (probe, destination IP) pair was used twice: once using ICMP traceroute and once using TCP traceroute. Using MaxMind, each IP address was geolocated at a country granularity, and with the resulting set of country-level paths, we analyzed which countries host and/or transit the traffic.

A client can use DNS open resolvers to discover georeplication of a site or service that can facilitate the avoidance of specific countries. For example, a client can query an open DNS resolver in a foreign country to discover a different georeplicated instance of a service; this technique (or the use of EDNS client subnet) can allow the client to discover different replicas of the same service. The path to this newly discovered replica may assist in avoiding particular countries (particularly if the client is trying to avoid the country of the original hosting replica!). The increasing use of IP anycast can sometimes make this technique insufficient, however, if the client receives the same IP address regardless of the apparent origin of the DNS query. Another approach is to use overlay network relays, which can prevent the client from traversing an unfavorable country by introducing a path that detours from the default. Additionally, using relays in the client's country can sometimes help keep local traffic local, by exploiting local paths that BGP does not select by default.

To evaluate the use of open resolvers as a tool for country avoidance, we query open resolvers in geographically diverse locations around the world using the Alexa Top 100 domains. Then we use RIPE Atlas probes to traceroute to the responses, and map them to country-level paths. To evaluate overlay network relays for country-avoidance, we establish 12 relays in geographically diverse locations around the world, run traceroutes from the country of interest to each relay, as well as from each relay to the Alexa Top 100 domains. After mapping the traceroutes to country-level paths, we measure which countries are avoidable.

Accurate IP geolocation is challenging, but our study does not require high-precision geolocation; we are interested in providing accurate lower bounds on detours at coarse granularity, and previous work has found that geolocation at a country-level granularity is more accurate than at finer granularity [3].

3 Preliminary Findings

Hosting diversity. We start by looking at hosting diversity, more specifically, how many countries a domain is hosted in. More diversity should provide for the potential to avoid more countries. About half of the Alexa Top 100 domains in the five countries studied are hosted in more than one country.

Routing diversity. Despite strong efforts made by some countries, their traffic still traverses surveillance states, and is subject to surveillance. Over 50% of the Alexa Top 100 domains in Brazil and India are hosted in the United States, and over 50% of the paths from the Netherlands to the Alexa Top 100 domains transits the United States. About half of Kenyan traffic traverses the United States and Great Britain.

On the feasibility of avoidance. By measuring which domains are accessible without traversing a given country using open resolvers and an overlay network, we see that there are ways to circumvent unfavorable countries. Without these country avoidance techniques, Brazilian traffic transited Spain, Italy, France, Great Britain, Argentina, Ireland (among others), but using the overlay network, Brazilian clients could completely avoid these countries for the top 100 domains. The overlay network can be used to keep local traffic local; the percentage of tromboning paths from the United States to the top 100 domains decreases from 11.2% to 1.3% when relays are used.

Cause for concern. Unfortunately, some of the more prominent surveillance states are also some of the *least* avoidable countries. Most countries are highly dependent on the United States, a known surveillance state, and not dependent on other countries. Neither Brazil, India, Kenya, or the Netherlands can completely avoid the United States with the country avoidance techniques. With the overlay network, both Brazilian and Netherlands paths avoid the United States about 65% of the time, and the United States is completely unavoidable for about 10% of the paths because it is the only country where the content is hosted. Kenyan traffic can only avoid the United States on about 40% of the paths from Kenya to the top 100 domains. On the other hand, the United States can avoid every other country except for France and the Netherlands, and even then they are avoidable for 99% of the top 100 domains.

Acknowledgments. This work was supported in part by NSF Award CNS-1540066.

References

- [1] Brazil Builds Internet Cable To Portugal To Avoid NSA Surveillance. <http://www.ibtimes.com/brazil-builds-internet-cable-portugal-avoid-nsa-surveillance-1717417>.
- [2] Netherlands New Proposal for Dagnet Surveillance Underway. <https://edri.org/netherlands-new-proposals-for-dagnet-surveillance-underway/>.
- [3] B. Huffaker, M. Fomenkov, and K. Claffy. Geocompare: a comparison of public and commercial geolocation databases. *Proc. NMMC*, pages 1–12, 2011.
- [4] J. Karlin, S. Forrest, and J. Rexford. Nation-state routing: Censorship, wiretapping, and bgp. *arXiv preprint arXiv:0903.3218*, 2009.